

METHOD AND APPARATUS FOR DISTRIBUTION OF FASHION AND
SEASONAL GOODS

Inventor: Robert E. Dvorak

BACKGROUND OF THE INVENTION

Allocation of goods among locations in a retail chain involves significant
5 stakes for operators, particularly for fashion or seasonal goods. A fashion program
typically consists of a number of items that will be sold for a fixed period of time.
For each item, the retailer typically purchases a fixed number of units in advance and
may not have the opportunity to reorder the item from the supplier. During the selling
season, the retailer needs to send product from its distribution center to each of its
10 locations to ensure that locations are in-stock in as many items as possible. As the
end of the selling period approaches, retailers typically reduce the price of unsold
items (markdowns) to ensure that all or substantially all of the inventory associated
with the program is sold at or near the end date of the program.

Determining the appropriate quantity of each item to send to each location
15 involves a number of considerations. It is useful to predict the rate of sales over any
period of time during a fashion or seasonal program, but this is difficult as sales rates
vary during the program. For example, in a holiday program, the sales rate in the
week before Christmas may substantially exceed the sales rate the week prior. The
relationship between sales rates in these two weeks may vary by product category,
20 price point, or other factors. Determining the characteristic shape of the selling

profile over time is difficult because of calendar differences between years and different promotional events from year to year.

Actual sales in particular locations vary widely from planned sales. Many locations may have small sales volumes, at least for specific items. Random variations in sales have a substantial impact on small sales volumes. Therefore, it is useful for an allocation algorithm to reflect the randomness of sales at specific locations, to avoid excess inventory in some locations and lost sales opportunities in other locations.

Retailers may designate presentation quantities for each location and item.

Presentation quantities are quantities adapted to a fixture or a location display goods to present attractively. It is useful for an allocation methodology to ensure that location presentation quantities are provided at the outset and maintained. Because retailers sell all of the product purchased for a program, location on-hand quantities eventually fall below the presentation quantity. The longer the location presentation quantity can be maintained, the more attractive the program appears to customers. In the case of coordinated products like a matching dress, blouse and vest keeping the presentation quantities for all of the products coordinated is also useful. It may be helpful to adjust and reduce presentation quantities as overall inventory is depleted or as a close out date approaches, to avoid excess stock in lower volume locations and lost sales opportunities at higher volume locations.

SUMMARY OF THE INVENTION

One embodiment practicing aspects of the present invention provides a computer-implemented method for adjusting a reference selling profile for a reference product, comprising retrieving one or more reference selling profiles corresponding to daily or more frequent historical data for one or more reference products, and adjusting the

reference selling profiles to correct for one or more promotions which impacted the historical data. Other embodiments and aspects of the present invention are described in below and set forth in the claims.

5

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates handling of a causal events which resulted in an artifact in historical data.

Figure 2 graphically illustrates handling of the stockout causal event which resulted in an artifact in historical data.

10

Figure 3 depicts historical corrections by assortment fill.

Figure 4 illustrates an interface for displaying the relative shares of locations or locations in product sales.

Figure 5 illustrates interface for displaying day-by-day diagnostic sales profiles, which sum to 1.0 over the selling period.

15

Figure 6 graphically depicts a share group.

Figure 7 illustrates an interface for entering or changing in season fractions or factors.

Figure 8 illustrates an interface for reporting a model stock of a particular item across all locations.

20

Figure 9 illustrates interface for reporting distribution of particular item across all locations.

Figures 10 and 11 illustrates an interface is for setting changing markdowns.

Figure 12 illustrates an interface for selection of markdowns.

Figures 13 and 14 graphically depict aspects of a markdown manager.

25

DETAILED DESCRIPTION

The following detailed description is made with reference to the figures.

Preferred embodiments are described to illustrate the present invention, not to limit its scope, which is defined by the claims. Those of ordinary skill in the art will recognize a variety of equivalent variations on the description that follows.

I. The Fashion Allocation System Overview

A fashion allocation system embodying aspects of the present invention forecasts sales and calculates allocation quantities for each distribution flow from a distribution center to locations in a fashion or seasonal program. While the fashion allocation system does not determine required vendor order quantities (also known as the total buy), the system can be used to generate information useful in determining the order quantity, especially if a product comes in many sizes, colors, flavors or other distinction where past location or other selling profiles are relevant. Ultimately, the total buy quantity decision is driven by a planning function and merchant assessments. After the decision on the total buy is complete, the fashion allocation system provides information useful in determining where and when to send the goods.

The fashion allocation system can take advantage of historical data for comparable items and then real sales as they become available. Historical data on a reference product helps forecast how the different locations will sell the different items. Using well-selected historical reference product(s) provides insights into sales quantities and sales timing by location on a relative basis. These insights can be adjusted based on actual in-season sales results. The fashion allocation system may be used to combine the historical and real sales results. A reference product may be an individual product, for instance a product comparable to a particular product for

which sales are being forecast, or it may be constructed based on user judgment or from a group of individual products.

In operating the fashion allocation system there are 5 activities, which involve several distinct aspects of the present invention:

- 5 1. Create Shares
2. Create Profiles
3. Setup Item and Location Data
4. Setup PQs and Causal Calendar
5. Run Distributions

- 10 The first 4 activities are typically done once for each fashion or seasonal program and then changed if required. The fifth activity, running distributions, is typically done multiple times over the season.

The fashion allocation system can use both historical reference product information and actual sales results in its forecasting, allocations and distributions.

- 15 The historical information is used in the form of shares and sales profiles, which are built for the new products using the reference products. Then item and location data and quantities (PQs) and a causal event calendar are set up. The casual event calendar can be used to capture promotional information such as advertising, preferential display of goods and price reductions or seasonal data such as the days preceding and
- 20 following Christmas, the days preceding and including Valentines Day and the days preceding the start of school (typically called "back to school"), which can be used to adjust either shares or profiles. With this information in place, distributions can be run. Once in-season sales are available, shares, profiles and distribution runs may reflect both projections and actual sales. This blending of projections and in-season
- 25 results is done in a way to not over react to random selling anomalies while not under

reacting to real differences in what was expected and what is actually happening in the new item selling.

5 A share or location distribution share is a set of location-by-location sales fractions adding to 1.0, where a locations' fraction is an estimate of its unit sales relative to the total unit sales of all locations. Shares typically, but not always, are calculated for groups of items instead of single items, so that the location level sales do not reflect individual anomalous situations such as random selling noise. Shares can be generated using a number of different historical sales periods (e.g., all the sales time for the historical items or just during a specified time period) and sales types
10 (e.g., actual sales or stock out filled sales). One share is typically assigned to an item, and many items may use the same share. Shares should differ for items that sell differently across the locations. For example, some locations may be strong sellers at high price points and therefore if the item in question is at a high price point, a high price point share may be used.

15 A reference selling profile is a historical set of data corresponding to day-by-day or more frequent sales for reference product(s). Typically, the historical sales data is stored as fractions adding to 1.0, where a daily sales fraction is the fraction of season unit sales that would occur on that day. Alternatively, a reference selling profile may store actual historical unit or dollar sales data, with appropriate scaling, as
20 will be apparent to those of ordinary skill in the art. However the data is stored, the invention can readily be adapted to historical data retained for intervals more frequent than daily, for instance, for three shifts a day. Like a share, a profile is typically determined for a number of past items so that the location level sales do not reflect individual anomalous situations. Unlike shares, profiles typically are grouped by
25 locations. In grouping locations, it is preferred to select a group sufficiently large to

overcome random selling noise but not so large as to blend together locations with different sales patterns. For locations that sell large quantities of particular items, grouping may be unnecessary..

The profile created may be adjusted to reflect to a new time period
5 corresponding to when the new items will be sold. So for example, a profile created for a first year during Thanksgiving and Christmas holiday may be shift to a second year reflecting the different days of the week and timing of events such as Thanksgiving. Likewise, a profile created in a first year in the spring may be time shifted for use in a second year during the summer. In making these types of changes
10 a reference selling profile may be truncated to a shorter selling period by dropping early or late sales dates and then rescaling the remaining days to more closely match the sales dates that are being projected rescaling the profile, if desired, to sum to 1. Alternatively, the reference profile may be compressed to a shorter period rather than dropping dates. For in-season selling situations profiles may be changed in a way that
15 they scale to a value greater or smaller than 1 reflecting whether unpromoted sales of a particular product should sell out a total buy quantity during the defined selling season (in cases where they will no markdowns are required or the season will take longer than planned). Profiles may be assigned at the item and location level or to groups of items or groups of locations. A number of items and locations may have
20 the same profile. Profiles may be adjusted to reflect timing difference among locations in product sales. For example, locations in areas like Florida or Arizona often have sales profiles influenced by seasonal influxes of people. Locations in some geographies may have different day of week and week of year variations in selling. Profiles for a specific location may also differ by price point or by type of goods.

008011 44680460

Profiles also may be adjusted to correct for seasonal effects which impacted sales of reference product(s). One seasonal effect is special selling days, such as the days before Mother's Day, Memorial Day weekend, the days before and after Thanksgiving, the days before and after Christmas. Depending on the product or location, any number of other special selling days may produce seasonal effects. Another seasonal effect is periods such as back to school days and Christmas shopping. Depending on the product or location, any number of other special periods may produce seasonal effects.

Corrections made to create shares can be applied to adjust reference selling profiles to generate projected sales profiles and vice-a-versa.

Once a share for an item and one or more profiles for locations have been selected, setup can proceed for specific items and locations. This may include setting the data elements that should remain the same throughout the distributions or change infrequently. These elements include the authorized locations for an item, the item buy quantity available for distribution over the life of the program and the in and out dates for the item at each specific location (if it differs by location). These may all be specified at the item or item/location level, as desired, and then are used by the system throughout the distributions. This data is typically modified if, as the selling progresses, if it becomes apparent that an incorrect share or profile was selected initially.

The presentation quality (PQ) and causal calendar are used to identify which promotional events will occur at each location and which items will be affected. Promotions, in this sense, may include advertising, preferential display or presentation of products, markdowns and the like. The PQs specify a desired quantity of inventory presentation in a location to be aesthetically pleasing. It can be specified on an item,

location, and day basis (i.e., need to have 12 on hand in location number 62 to display from December 5th to December 12th.) The causal calendar identifies promotional events (again, on a location, item, date basis) that will affect each item during the life of the fashion or seasonal program (i.e., item number 876 will be on 25% price reduction from December 1st through December 7th at these specified locations). The causal calendar will also identify the expected sales lift associated with the event (e.g., a 2X multiple of normal selling) and may use a value that is user a specified or historically derived.

Distributions can be run and reports generated after setup. The distributions suggest actual allocation quantities for each location, given the appropriate profile, share, actual selling, presentation quantity requirements, causal calendar, and total buy for the product. These distribution results are typically uploaded to another system for distribution center picking and shipping to locations.

II. Create Shares

One aspect of the fashion allocation system is an historical share. To recap, a share is an historical set of location-by-location sales fractions adding to 1.0, where a locations' fraction is an estimate of its unit sales relative to the total unit sales, typically over the total season. A single share typically, but not always, is created using a group of items, so that the location level sales do not reflect individual anomalous situations such as random selling noise.

A share is created using historical sales data, which may previously have been loaded into the fashion allocation system. It is useful to select a set of past items with selling drivers similar to the selling drivers for the new items for which the shares being created. The right time period to use and type of sales (actual or filled) also may be taken into account when selecting historical data with which to create a share.

Selection of raw historical data is accompanied by selection of desired historical corrections. This includes selecting a type of stock out or out of stock correction method to use and a type of filling. Other corrections may include correction for promotions, including corrections on a location-by-location basis, and
5 correction for seasonal effects. After the selections are made, the fashion allocation system can generate a share to which the user may assign a unique name.

A share is typically, but not always, built up from the historical sales of a number of historical items or reference selling profiles. For each of those items it uses the actual unit sales, on hand inventory (if selected and available for determining
10 stock outs) and information from the causal calendar. In building the share, the fashion allocation system may remove the past impacts reflected in the causal calendar so that different promotions done across different locations (if it was done that way) or different impacts of seasonality do not influence the calculated shares. Such promotions may include advertising, preferential display or reduced pricing of
15 reference product(s). Such seasonality may include different holiday periods such as Christmas, Easter and Mothers Day and other periods like "Back to School". Once a share has been constructed, it can be automatically adjusted to reflect future promotions or seasonality entered into the system on a location-by-location or higher-level basis (e.g., regional or total item).

20 The fashion allocation system enables user selection of past items that are expected to sell in the same manner as new items. The user's thinking should take into account a number of factors such as target consumers, price point, fashion ability, branding, and size selling. In addition, it is some times necessary to think about factors such as male to female crossover sales or cross assortment impacts. One

usually does not need to think about promotional impacts, as one aspect of the fashion allocation system is adjustment for promotions.

The group of items used to create a share is selected. This selection depends in part on the sales volume of the items being selected; fewer items are required when the sales volume is large. One trade-off worth considering is selecting enough items to avoid a profile that is driven by a few anomalous events, but selecting a few enough items to avoid averaging away discernible selling differences. That is somewhat driven by the unit sales volume for the past items. For example, if particular styles or size splits have relatively predictable location selling differences, it is preferred to make sure that the shares created have those splits and not some higher level grouping that averages away those impacts.

The fashion allocation system allows the user to test shares after they are created, by running different alternative groupings and then evaluating whether those groupings show predictable differences over time. Typically, groupings are refined over time and change as consumer preferences and trends change. One strategy is to start with somewhat larger groupings and develop smaller groupings based on experience. Beginning with small groupings exposes the user to differences that are potentially anomalous and not replicable. Choosing a broader grouping takes advantage of a function of the fashion allocation system that corrects differences between projected and real sales, as real sales results become available.

After selecting a reasonable set of comparable products, a time period should be selected to use for the share that is being created. It is useful to select the time period with the most representative sales. The fashion allocation system includes an option to make this decision less critical by using a sales type that fills stockouts. The system may compensate for stockouts throughout and at the beginning and end of the

sales period, to avoid having a particular location underrepresented in the share calculation. Alternatively, a time period that ends before the best selling locations in the chain run out of the comparable products can be selected, or a combination of the two approaches.

5 Corrections may be made when historical data reflects stockouts of items. The number of historical corrections to be made depends on the quality of the data selected. Software can interact with a user in determining the stockout types present in the data and offers filling options. The user may decide whether to accept the default or to alter the default assortment fill weighting method. Alternatively, the
10 software may select the fill method to be used.

Logic for correcting data by filling stockouts is illustrated in figures 1 and 2. The system may provide for overriding any of the stockout types using straight or orphan logic. Both straight and orphan logic determine when sales start and stop (depending upon user specified parameters) and, as necessary, override in stock
15 assessments. In figures 1A and 1B, the data used by the system includes calendar date 101, whether or not stock was on hand 102, the sales on the specific date 103, the sales program starting date 104, and the sales program ending date 105. At the beginning of the program, for early sales, using the straight sales logic the system interprets the lack of sales 106 as a stockout condition. Although the on hand data
20 102 suggest that inventory was available for sale, the sales pattern suggests that the inventory was not on display. The system may present the user with options for filling the stockout period from October 1 to October 4 , 1998 (labeled 106).

At the end of a sales program, the system interprets the lack of sales on January 5, 2000, labeled 107, consistent with the on hand data 102 as indicating that
25 the inventory stockout immediately followed the program close out date 105. In this

example, a single sale 107 nearly two weeks after the earliest prior sale may be flagged by the system as a likely mistake in the data, which the user may choose to override. A threshold can be set to identify orphan sales which should be ignored for purposes of detecting stockouts at the beginning or end of a sales program, such as one or two units sold at a location separated from other sales by two or more days of no sales, at the beginning or end of the sales program.

Two different options for filling stockouts, average and profile filling, are graphically illustrated in figure 2. The average filling approach calculates the average sales for in stock days during each promotional period and then applies the calculated average to each day stocked out. The profile filling approach uses the stocked out day's percentage share, from a profile, to calculate the filled sales. If a user elects to use profile filling and the item is not associated with a profile, then a profile may be assigned for use or average filling can be used. Other forms of filling or other forms of average filling may be used, such as taking the sales for a number of days around the stockout day and using the average sales for those days.

An additional historical correction may apply when shares are based on an assortment of items, particularly when some of the locations did not stock all of the same products as other locations. This correction is called an assortment fill.

Assortment fill applies to situations where the item was not stocked at a location; this is different from stockout filling which applies to individual days when an item stocked at a location was out of stock. Figure 3 illustrates this. Cell 301 is blank because product 105423 was not stocked at some particular location. One option is to fill the cell, so that the total shares by location calculated and displayed in column 302 are not prejudiced by the assortment of goods carried at the particular location or by the group of products selected as a basis for calculating shares. The assortment fill

function provides that locations that did not sell one or more of the items used to calculate a share will not be given an inappropriately low share.

The type of assortment fill weighting used may be selected when a user creates shares or profiles. Two different types of weighting include: unit sales and percent sales. Unit sales weighting is based on total units sold and is useful for large unit volume items. Percent sales weighting is based on the share for each item. Other forms of weighting may be used such as the dollar sales of each of the items. Selecting the best form of weighting for a situation can vary, for example if for a high volume item units may be the best approach while for a high price item dollars may be the best and for a typical item percent sales, but is important to generating a most representative weighting of the items.

While stockout corrections are presented in this example in the context of adjusting shares, the same method applies to adjusting reference profiles and to projections made from historical data impacted by stockouts.

Selection of the items of interest, the time period, the stockout type, the stockout filling, and the assortment filling selected primes the system to create shares. The shares algorithm typically is run for all of the locations for which an operator is responsible. Being over inclusive of locations does not distort the relative shares of a subset of the locations; being under inclusive requires that shares to be recalculated. Calculated shares are named. Figure 4 is an example of calculated shares. Column 401 specifies location codes. Column 402 give names of locations. Column 403 gives unit fractions for specific locations. These unit fractions total 1.0.

III. Create Profiles

In the example described here, shares are created before profiles. However, these aspects of the present invention may be practiced separately or in a different order. In the course of setting up shares, most of the data has been collected to create
5 profiles. The required data and resulting outputs are described in a table below. A profile is a set of historical data corresponding to day-by-day or more frequent sales. Typically, this set of data is stored as fractions adding to 1.0, where a daily sales fraction is the fraction of season unit sales that would occur on that day.

Alternatively, actual sales volume can be stored, subject to scaling as necessary.

10 Again, however the data is stored, the invention can readily be adapted to historical data retained for intervals more frequent than daily, for instance, for three shifts a day. Like shares, a profile may be determined for a group of items so that the location level sales do not reflect individual anomalous situations. A user typically selects locations to group together which could, but does not have to, be all locations. Those
15 groups should have similar time selling profiles and the group should be sufficiently large to overcome random selling noise but not so large as to blend together locations with different time selling.

A system embodying aspects of the present invention may provide utilities to generate a diagnostic profile output for an item or set of items and a number or all of
20 the location locations. This is done by selecting the stockout type(s) and stockout fill. In addition, a share/profile weighting needs is selected. The diagnostic profile output provides a day-by-day unit sales profile for each of the locations specified, enabling a user to look for unexpected differences. Unexpected differences can be evaluated to determine whether they are legitimate differences or result from anomalies.
25 Diagnostic profiling facilitates grouping locations that have roughly similar profiles.

Profiles can be run by group. Users may prefer to have fewer rather than more profiles, to evaluate the differences among profiles, and to make judgments about whether differences observed in the past will continue into the future. A sample output for a uniquely named profile is shown in figure 5. Column 501 lists dates
5 covered by the profile. Column 502 is unit fractions for a profiled period, which add up to 1.0.

Profiles can be created based on more than one reference product or more than one reference profile. It may be useful to create a seasonality selling profile, distinct from the product selling profile. A general profile reflecting seasonal selling can be
10 created based on sales of a group of more basic or non-seasonal items. Alternatively, it could be created based on sales of a group of non-promoted seasonal items. A reference selling profile can then divided by or ratioed to the general profile to give a seasonally adjusted product selling profile. A seasonally adjusted selling profile can usefully be applied to a different calendar period as a reference profile. Otherwise,
15 days can be shifted to reflect differences in the calendar from year to year and time period to time period.

In addition to the creation of shares and profiles, a user typically works with the fashion allocation system to set up items and locations. However, the order in which the user carries out the steps described in this example is optional; many of the
20 steps can be carried out in no particular order or in parallel.

The set up of items and locations for the fashion allocation system, typically is done once. Set up information should be modified when the sales program changes and when there are refinements to what is done. There are a number of settings that are done at the item level and that are done at the item/location level.

INPUTS (OUTPUTS)		
Input	Description	Considerations
Item level		
Share	A share	Select an appropriate named share.
Share Group	Set of items used to make in-season share adjustments	Which items have similar selling characteristics
Available Buy	Quantity available for specified locations	The total buy for these locations
DC availability override	Quantity available for distribution	What will really be available
PQ distro minimum	Flag to specify whether to use PQ minimum or not	Whether you want each location to get a minimum of the PQ
Case distro minimum	Flag to specify whether to use case minimum or not	Whether you want each location to get a minimum of one case
Coverage index	Safety stock-like coverage factor varying by sales percentage	Balancing safety stock vs. liability of putting stock in wrong place
DC availability override	Override quantity for distribution in units	When you want to run a distro, distributing more stock than is in the DC
Location level		
Profile	A profile	Select an appropriate profile for an item/location
In date	Date when item is ready for sale at location	In-location setup date
Out date	Date when item should be sold out	Plan
Pivot date	Date when specified item sell through desired	Whether there is some special timing driving the use of a pivot date
Pivot before percent	The percent sell through specified by the pivot date	Plan

Following settings may apply at the item level. Because the same settings often apply for a group of items, the fashion allocation system may facilitate making settings for an entire group. For convenience, it may allow a user to make settings for

5 a group and then go back and change the settings for exceptional items.

A user may select the item or items that they want to set up and specify a share to use, a share group to use, whether a PQ and/or purchase minimum is required, a coverage index to use, and whether a distribution center availability override is desired. Specifying the share typically involves selecting the name of a previously created share to be used. A share group is a set of items that are grouped together for in-season share determination purposes, for example the share group could be a single item, a single style, many items of the same size across a number of styles or a number of items or styles that share common characteristics. As with a share, specifying a share group involves selecting the name of a previously created share group.

A share group serves an aggregation function similar to that of the share, constructed for group of items. Figure 6 illustrates a hierarchical share group. Column 601 is a list of items in an overall share group. Column 602 symbolizes the highest categorical level of the share group. Column 603 subdivides column 602, as into premades and single items. Column 604 further subdivides part of column 603, separating premades into branded premades and other premades. Finally, column 605 further subdivides the column 604 branded premades into premades from Brand A, Brand B and Brand C.

Once in-season sales began, the fashion allocation system may be capable of modifying the share used in creating the model stocking level based on actual sales. This modification is applied by location. However, in cases where the selling quantities are relatively small at the location level, the system may avoid modification to the model stock based on a single item/location sales results. This may be avoided because the usual randomness of selling results in some locations selling a number of items quickly and others selling none; both of these events are within the expected

range and variation of selling. Therefore, to quickly pick up real trends in selling, the system may to use a share group to look at a number of items selling in a location and make adjustments based on a share group, rather than a single item.

Upon completion of initial data input, a system practicing aspects of the present invention may prompt the user to specify for each location whether the location will receive a minimum of one case of product, the economic stocking level, the item/location specified PQ or the larger of the economic stocking level or specified PQ. This assumes that there enough units available to satisfy a selected minimum. However, if the minimum is not satisfied, the system should do its best to fulfill the requirement and then let the user know that there was insufficient available product. Then the system may attempt to fulfill the missing quantity next time that a distribution is run and continue to do so until the minimum is fulfilled.

A Coverage Index can be specified. The Coverage Index specifies a multiple of expected sales for the coverage cycle that will be distributed in the first allocation to the locations. It is similar to so called safety stock for basic goods; however, it does not function quite like safety stock, because all of the fashion goods in a sales program are intended to be distributed in a specific time period. Any so-called safety stock sent out early in the sales program serves as a cushion against sales occurring faster than expected, but also amounts to a potential liability for distributing stock to a slow selling location.

Figure 7 illustrates an interface for entering or modifying in season fractions or factors. As actual sales results become available, the mix between relying on projections and actual sales may shift. In season fractions may automate the process of shifting this mix. Figure 7 illustrates five different weighting factors applicable in five different in season periods. The in season periods may correspond to fractions of

The fashion allocation system process also involves specifying a distribution center availability override, if applicable. A key distribution center override allows the user to run a distribution with a quantity greater what is available in the default distribution center.

Location setup can be done for a single item/location combination or it can be done for many combinations of the same time. If it is done for many, each of the many will get the same setup. In some instances, the quickest way to input certain settings may be to specify all of them at once and then edit the few exceptional items/locations. The user may proceed to set up and specify the profile to be used, a sales in date, a desired sales out date or close out date and, if the sales program includes it, a pivot date and pivot before fraction. Specifying the profile involves associating particular item(s) and location(s) with a profile. Specifying the sales in date is a matter of inputting the date when an item is expected to be in the location ready for sale. The out date is the date by which the item preferably is completely sold out this location. Finally, a pivot date is used in special event selling situations where the sales plan involves having sold a specified percentage of the goods by a specific date. Specifying a pivot date anchors a sales plan, so that the specified percentage of goods to be sold by the pivot date remains the same, even if the in or out dates change. For example, if it is preferred to have sold 90% of units of a

product by Christmas day, then the Pivot Date would be December 24th and the Pivot Before Fraction would be 0.9.

IV. Setup PQs and Causal Calendar

5 Presentation Quantities (PQs) may be set for each item, location, and time period. These PQs represent potential increases to the allocation quantities that would otherwise be calculated. The PQs are typically set to ensure that fixtures within locations are visually appealing with sufficient product on hand, presumably increasing customer propensity to buy. These PQs may be subject to a cap. Such a
10 cap may be a proportion of the projected remaining sales, for the life of the sales program. Or, it may be a factor of the ratio of stock on hand at a location to projected remaining sales. Over time, such a cap naturally may be reduced, much as the reliance on actual sales as opposed to projected sales is increased over time. It is useful to reduce and eliminate the use of PQs for distribution as the end of a sale
15 program approaches because stock on hand must drop below the PQs in order to sell out or close out the stock.

 The causal calendar can identify historical or planned promotional events on an item, location, and time period basis, with the user analyzing sales lifts associated with historical promotional events or setting the expected sales lift associated with a
20 planned event. For example, an item may be promoted with a 25% price discount from one date to another date in all locations, and this might result in an expect 100% increase in sales. These causal calendar entries will be used to adjust the projected sales profile at the location-item level. These adjustments can be made by historically based sales lift calculations, current in season selling lift calculations, a blend of

historical and current calculations or different variants of specified values.

V. Run Distributions

- The table below identifies some of the inputs and outputs that may be useful
- 5 for running a distribution.

ELECTRONIC INPUTS	
Sources	Data elements
Other elements	<ul style="list-style-type: none"> • Shares • Profiles • Item and location setup • PQ and Casual Calendar • Actual Results
Shipments	<ul style="list-style-type: none"> • Item/location shipments so far
Shipment changes	<ul style="list-style-type: none"> • Any discrepancies, transfers or other impact to item/location shipments
Actual	<ul style="list-style-type: none"> • Actual unit sales for the item/location so far
USER INPUTS	
Inputs	
Distribution Date	When distribution will take place
First Distribution Flag	First distribution
Last Distribution Flag	Last distribution
Next Distribution Date	Date of next distribution (if this is not the last distribution)
OUTPUTS	
Element	Description
Item/location distribution	Quantity of goods to be sent to each location

As the season progresses, the fashion allocation system may automatically react to actual sales and modify the distributions to reflect overall item and individual

location sales trends. The fashion allocation system can be programmed to highlight major differences. Major differences may then be evaluated to determine whether the differences flagged indicate are a real trend that should be reacted to, or whether there are some extenuating circumstances. The system logic and resources can react automatically to trends; this automatic reaction may be tempered or overridden by a user when there are extenuating circumstances.

Using shares and profiles, PQ and causal calendar information, a system practicing aspects of the present invention can easily produce a distribution. To do so, the user sets up and names an activity. That activity specifies the items and locations to be included in the distribution. That activity is "run" and then saved for reuse. To run a saved activity, the user may simply specify the date of the intended distribution, whether that distribution is a first, last or intermediate distribution and if it is not the last distribution then the date of the next distribution. With these or a subset of these parameters, the distribution run can proceed.

The user specifies the reports to be printed for the distribution. Two typical reports include the Model Stock report (figure 8) and the Distribution report (figure 9). The Model Stock report includes key inputs and outputs for determining the item/location cumulative model stocks. The Distribution report then starts with the model stock and works through to the actual quantity for distribution to each location. An approved Distribution report may then be uploaded for distribution by a Distribution Center.

Figure 8 shows a layout of the model stock calculated for multiple locations for a single item. The column for Location Number is the number assigned to the location by the retailer; Location Description is the description assigned to the location by the retailer; In Date is the first selling date of the item in the location; Out

Dates may be the last intended selling date or the close out date of the item in the location; Activity Cycle End is the expected receipt date in the location of the distribution subsequent to the one for which the model stock is being calculated, or, alternatively, the out date if no more distributions are planned; Profile Name is the name that has been assigned by the retailer to the specific profile being used for the purpose of distribution; PQ is the presentation quantity assigned by the retailer for this item and location in order to ensure an aesthetically pleasing presentation; Plan Share Percent is the share for the location and item calculated in advance of the season.

Group Sales Share Percent: the share for the location and item based on actual selling in-season; Interweave Share Percent is the Plan Share adjusted by the Sales Ratio (as defined in the distribution calculation described below.) This represents a weighting between the Plan Share and the Group Sales Share Percent depending on the Interweave Factor in effect. Cumulative Model Stock is the cumulative model stock for the item and location. Sales Thus far is sales to date for the item and location.

Figure 9 is a Distribution Report that shows the actual distribution quantities calculated for location for a given item. In the report Location Number is the number assigned to the location by the retailer. Location Description is the description assigned to the location by the retailer. In Date is the first selling date of the item in the location. Out Date may be the last intended selling date or the close out duty of the item in the location. Activity Cycle End is the expected receipt date in the location of the distribution *subsequent* to the one for which the model stock is being calculated. Cumulative Model Stock is the cumulative model stock for the item and location. Previous Shipment is the quantity of the item that has already been shipped to this location. Desired Shipment is the desired shipment quantity to the location (equal to cumulative model stock less previous shipment). Case Rounded Shipment

(Units) is the desired shipment constrained by rounding to the nearest integer case quantity or to the next larger integer case quantity but constrained by the total available to distribute. Case Rounded Shipment (Cases) is desired shipment as above, in cases as opposed to units. PQ Distro Minimum is user input flag, where Y indicates that any calculated shipment quantity will be overridden to ensure the distribution of a minimum of the presentation quantity for the location, item, and time, and N indicates that it will not be overridden. Case Distro Minimum is user input flag, where Y indicates that any calculated shipment quantity will be overridden to ensure the distribution of at least one case per location and N indicates that it will not be overridden. Insufficient PQ is a calculated flag, that indicates if the system was unable to satisfy the presentation quantity needs in all locations.

The actual distribution quantity can be calculated to take into account the model stock, previous shipments to the location, such factors as transfers, sales and other factors impacting quantities available such as damaged and therefore non saleable items. The following formula is useful in calculating model stock:

$$ModelStock = TotalBuy_{item} \bullet (Share_{item, location}) \bullet (Pfile_{item, location, date1, date2}) \bullet (SalesRatio_{item, location}) \bullet (CoverageIndex \% sellthu) + (CappedPQ_{item, location})$$

In this equation, $TotalBuy_{item}$ is equal to the total quantity of the item available to be sold during the entire length of the program. It may also be called the "buy quantity" or "total buy." $Share_{item, location}$ is equal to the share calculated for the specific item and location as described earlier. It represents the single location's expected selling quantity of the product as a fraction the total. $Pfile_{item, location, date1, date2}$ is equal to the sum of the daily profile quantities for the item and location from day1 to day2, where day1 is the date of the start of the item selling, frequently called the in date, and day2 is the receive date of the next distribution. (This formula and the invention can

readily be adapted to historical data retained for intervals more frequent than daily.)

For example, if the in date were November 1, the next distribution date were

November 8, and the shipment would take five days to get to the location, then the

Pfile would be equal to the sum of all the daily profiles from November 1 to

5 November 13.

SalesRatio_{item,location} is equal to (*UnadjustedSalesRatio_{item,location}*) • (1-*factor*) +

(1•*factor*), where the *UnadjustedSalesRatio_{item,location}* is equal to the ratio of actual

selling to projected sales. In cases where the sales ratio is calculated to be less than

1.0 the system, depending upon user specification or automatic set up, a user

10 optionally may elect to use 1.0 instead of the smaller number. The rationale for that

option would be that the selling is below plan but markdowns can and may be made

on short notice to increase sales to the plan level and therefore distributions should

anticipate that possibility and have sufficient goods out in the locations. When a

distribution is run, the projected sales for a time period typically equals *TotalBuy*

15 multiplied by the sum of daily profiles (fractions adding to 1.0) from the beginning of

the program through the day prior to the distribution run. A *factor*, such as the in

season factor, may be user specified as a function of sell-through or elapsed part of

the planned selling period. At low sell-through levels, it is useful for the *factor* to

approach 1.0, while at higher sell-through levels, it may be more useful for the *factor*

20 to approach zero. This *factor* determines the relative weighting of planned sales to

actual sales.

In the formula above, *Coverage Index* is an index typically specified by the

user during user setup. During setup, the user can specify a range of coverage indexes

that vary over time, for instance linked with the percent sell-through of the product.

25 Typically, it is useful to start with higher values (e.g., 3) and then drop to 1 as the

product approaches complete sell-through. $CappedPQ_{item,location}$ is equal to the desired presentation quantity set for the item and location for the time periods relevant for the distribution, but reduced such that the sum of the presentation quantity and the balance of the calculated distribution quantity does not exceed a fraction of remaining expected sales for the item and location for the balance of the program (as calculated by $Total\ Buy \bullet Share \bullet Profile \bullet SalesRatio$).

A system practicing aspects of the present invention can automatically react to in-season selling results by altering the distribution of goods to match sales trends.

For instance, when selling starts and the item moves faster than plan, the fashion allocation system can immediately adapt to the higher selling sales rate and increase the size of the next distribution to catch up. Since the fashion allocation system has a safety stock-like cushion built into its early distributions, usually an adaptation to a higher than projected sales rate can take place without sales disruptions.

The fashion allocation system also can react automatically to any changes in the causal calendar. When a user adds, subtracts or changes promotions, the impact on sales can be automatically reflected in subsequent distributions. Likewise, the fashion allocation system can automatically react to markdowns. In typical cases, a markdown is taken because sales have not met expectations, leaving more inventory in the locations than desired. Therefore, the markdown may be a "catch up" to move actual sales closer to planned sales. However, if it results in a faster than planned sales rate, the fashion allocation system can react at the next distribution and increase the flow of goods to compensate for the markdown actions. It also can alert the user to adjust additional scheduled markdowns.

A system practicing aspects of the present invention may be able to react to trends in in-season selling in two or more ways. First, if the Share Group Factor for

an item and location is changing, the distributions can be increased or decreased to reflect that impact on the item/location Model Stock. That Share Group Factor will change if the items that constitute the Share Group are selling above or below plan. That trend in sales can be reflected in the calculation of Model Stocks. There is also a second mode applicable at the item and location level. If an item has sold more or less than planned to date, the specific difference can be added to or subtracted from its next distribution. In other words, if an item/location has sold 10 more than its plan to date then it will be shipped an extra 10 in its next shipment to replace those sold. However, this replacement should not necessarily be on a one-to-one basis, as such a reaction might result in an inordinate increase in model stock for that item and location. Response to changes in the Group Share Factor, instead of an individual items sales, can reduce over reactions to random sales events.

The system's capacity for automatic in-season changes makes it useful for users to exercise judgment and to make sure that anomalous events do not precipitate unwanted, automatic changes.

Monitoring of in-season selling vs. planned selling is one way to test whether the user selected the right shares and profiles. Two indicators reported by the system are whether there are large differences in the plan versus actual location sales shares (usually measured for the share groups) and whether the week-by-week sales rate (relative to plan at the item level) is moving a great deal. If the Share Group Factor is moving a great deal for many if not all of the locations then the user may not have used the right share. The user may have selected a special event share when they should have used a more regular selling share or visa-a-versa. If that is the case, the system can easily be used to correct the error and rerun the distribution.

If the Share Group Factor is moving a great deal for a single location or a small number of locations, then the user may be alerted to check whether there has been some abnormal selling impact at those locations. Goods may have been delayed or set up early. Locations may have under or over presented the items. The user also
5 may consider whether the right share was selected.

In addition, distributions may encounter shortages in product available to be distributed at the time desired. This may be during the middle of a sales program, while awaiting additional inventory, or at the end of the program. In the instance of shortages during the program, logic may be provided to best decide where to send the
10 limited amount of goods. One logic provides for proportional shipping to each of the selling locations. Other logic may take into account case quantities the item(s), so that full cases are shipped, or for other user definable priorities such as filling the presentation quantities first. These factors can automatically be factored into the actual distribution quantity sent to each selling location.

VI. Markdown Manager Overview

The Markdown Manager includes aspects of the present invention which can assist the user in determining the timing and magnitude of markdowns of fashion and
20 seasonal items, as they impact gross margin, profit, sell out date, remaining on-hand inventory or other measures of effectiveness.

Based on the sales profiles associated with items and locations during the fashion allocation system setup, Markdown Manager can create an array of potential markdown scenarios, based on the allowed markdown levels and markdown dates,
25 and calculate revenue and other measures of effectiveness all or some of these

scenarios. The scenarios are run for combinations of markdown levels and dates.

These combinations can be specified by a user or generated automatically from lists of markdown levels and dates. Rules can be applied to the markdown levels, such as markdowns should become more attractive during the selling season, not less

5 attractive. One or more scenarios may be reported. Typically, the markdown report produced by the Markdown Manager is then reviewed by merchants and planners to accept or adjust the specific system recommendations, and then the resulting markdown decisions are entered in the causal calendar.

Five activities comprise the Markdown Manager operation:

- 10
1. Select items and locations
 2. Select markdown dates
 3. Setup markdown discount percentages
 4. Run Markdown Manager
 5. Enter selected markdowns into the causal calendar

15 In addition to the set up described above, item/locations may be associated with to pivot dates, which anchor markdown activities even if the in and out dates for a sales program are modified.

Setting up a markdown activity includes selecting items and locations for which markdown scenarios will be generated and reviewed. An interface for doing
20 this is depicted in figure 10. One option is for the Markdown Manager to generate scenarios using the same markdown for each product across all of the different locations; that is, so pricing is always constant from location to location. One input is a specification of the items and locations for running markdowns scenarios.

In many cases, a useful way to create a markdown activity may be to copy a
25 fashion allocation system activity that was created for fashion distributions. If all of

the fashion activity items are eligible for markdown at all of the locations to which the items were distributed, copying the existing activity may be substantially all that is required. If select items are not eligible for markdown, they may be removed from the markdown activity before or after copying. Even items ineligible for markdown
5 can be included in a run in order to see what markdown recommendations emerge, because system recommendations can always be ignored.

Setting up a markdown activity also includes specifying the eligible markdown dates. An interface for doing this is depicted in figure 11. These dates indicate to the system which days markdown actually may occur. The system can test
10 all eligible markdown levels at each of these dates to generate potential markdown scenarios.

The eligible markdown dates should begin after the start date of the program and the last date should be before the conclusion of the program. For a Christmas program, for example, eligible markdown dates might include weekly Wednesday
15 markdowns beginning the Wednesday before Thanksgiving and ending the Wednesday before Christmas. In this example, each of the Wednesday dates would be entered into the system. The system could then use those dates as potential dates to begin a markdown.

Preparing for a markdown run further includes associating items with potential
20 markdown levels and associated sales lifts. An interface for doing this is depicted in figure 12. If the same potential markdown levels (e.g., -25% and -50%) apply to all products and the sales lift associated with each of these markdown levels is identical across all items, then all that is required may be entering and selecting these two price levels and then applying them to the entire product tree. If different potential
25 markdown levels apply to different departments or even different items within the

same department, then all of the markdown levels may be entered into the system and associated with the items. Alternatively, sales lift data can be stored in a library or database accessible to the user, so that a sales lift can be selected based on past experience.

5 The markdown levels specified for an item are the levels tested. For example, if the eligible markdown levels are -25% and -50%, then these are the levels that they system will recommend. In providing the recommendation, the system will determine the appropriate date (constrained by the set of dates specified in the activity) that the markdown should take place or if a markdown would be counter productive.

10 In one embodiment of the present invention, each markdown level is assigned a descriptive name (such as "25 off"), a discount percentage (0.75, in this case, to represent 25% off), and a lift percentage. The lift percentage indicates the sales acceleration generated by a markdown of this level. For example, a lift factor of 2.0 associated with a discount percentage of 0.75 indicates that a 25% off discount results
15 in a sales multiple of 2.0. If different items respond differently to the same discount percentage, the same discount percentage can be entered multiple times in the table. For example, "25% off fragrance" and "25% off colorings" will both have a discount percentage of 0.75, but they may have different lift percentages. When matching
20 markdown levels with items, the appropriate markdown discount and consequent lift should be selected.

 A further aspect of the present invention is using the system to analyze and evaluate past sales results to determine the markdown lift percentages. This analysis can take the form of a graphical comparison of projections based on early actual sales versus historical actual sales after a markdown, or a more analytical lift percentage or
25 factor.

With the setup activities complete, the Markdown Manager can be run.

Typically, the Markdown Manager will be run one to three days in advance of each of the potential markdown dates, so that the results can be reviewed, the markdowns decisions made, and the decisions implemented on the markdown date.

5 Figures 13 and 14 illustrate how the Markdown Manager evaluates alternative scenarios. It first generates all valid scenarios of markdowns, based on all permutations of allowed markdown prices and allowed markdown dates. Scenarios where the price increases over time (caused by a larger markdown followed by a smaller markdown) may be excluded by a rule. Markdown Manager then calculates
10 total revenue and the number of clearance units for each of the scenarios generated. One or more of the highest revenue scenario may be selected for inclusion in the markdown report. Other measures of effectiveness can be calculated and used as a criteria for inclusion in a report. Figure 13 depicts how the markdown manager detects a short fall of actual sales, as compared to the *a priori* selling profile. Actual
15 sales are represented by the line 1301. Actual sales information is available through today's date 1302. Potential markdown dates 1303, such as Wednesdays during the Christmas selling season, are entered into the system. The original selling profile 1304 is the reference for comparison against revised projected sales 1305 which can be adjusted to reflect actual sales through today's date. Figure 14 depicts alternative
20 sales profiles based on an array of alternative markdown strategies. The markdown dates are 1403A-1403C. The end of program "out date" by which all units are supposed to be sold is 1410. The scenarios in this figure are no markdown 1405, a 25 percent markdown 1406, a 25 percent markdown followed by a 50 percent markdown 1407, and a 50 percent markdown 1408. From the curves, it can be seen that much
25 less than the entire stock 1409 will sell out if the no markdown strategy 1405 is

followed. The entire stock will sell out early if a 50 percent markdown 1408 is offered early. Other strategies produce intermediate results. A report generated by the markdown manager quantifies the results shown graphically in these figures.

A markdown report on paper, displayed at a terminal or otherwise presented,
5 can identify, for each item, the recommended markdown level at each of the specified markdown dates, a measure of profitability, a sell out date and a number of units left on hand after the close out date.

After the markdowns are reviewed and markdown decisions are made, these decisions can be implemented by entering them into the causal calendar so that future
10 distributions and markdown analysis will reflect the decisions. For each markdown decision made, inputs should be made to reflect the item, the markdown level, and the effective dates. Once the markdown level and time period are entered into the causal calendar, the system will be able to appropriately calibrate selling performance for the next markdown and make the appropriate distribution recommendations.

15 While the present invention is disclosed by reference to the preferred embodiments and examples detailed above, it is understood that these examples are intended in an illustrative rather than in a limiting sense. It is contemplated that modifications and combinations will readily occur to those skilled in the art, which modifications and combinations will be within the spirit of the invention and the
20 scope of the following claims. Each method may readily be recast as a device or article of manufacture.

We claim as follows: